

## BATCH - Bug # 1323: why "atmos" jobs beat up the database

<b>Status:</b>	Resolved	<b>Priority:</b>	High
<b>Author:</b>	Art Kreymer	<b>Category:</b>	
<b>Created:</b>	05/25/2011	<b>Assigned to:</b>	Robert Hatcher
<b>Updated:</b>	05/27/2011	<b>Due date:</b>	
<b>Subject:</b>	why "atmos" jobs beat up the database		

**Description:**

Date: Tue, 24 May 2011 16:55:53 -0500  
From: Robert Hatcher <rhatcher@fnal.gov>  
To: Minos Mailing List <MINOS\_SOFTWARE\_DISCUSSION@listserv.fnal.gov>  
Cc: Andy Blake <blake@hep.phy.cam.ac.uk>  
Subject: MINOS: why "atmos" jobs beat up the database

In today's Batch meeting I said I'd look further into why the database was getting hammered on the BEAMMONSPILL table. Art, Rashid and I ruled out that this is somehow caused by the keep-up jobs: the queries are on the dogwood1 DB and the problem extends long beyond when the last of the morning's keep-up files was complete. So it has to be the ongoing "atmos" jobs.

I re-tested BDSPillAccessor and for a day's worth of sampling at 1 second intervals it pulled only ~26 BEAMMONSPILL entries (one every 2 hours + once for each 1 second gap between regular entries). This is as I walked from T=2008-10-20 + 0 sec to T'=2008-10-21 + 0 sec at querying 1 second intervals in order.

I had Rashid increase the verbosity of DBI and run one of the "atmos" jobs that ran this morning during the time in question. What we found was a flip flopping of SEQNO's, that is SEQNO\_1, backstop, SEQNO\_1, backstop, SEQNO\_1, backstop, ... rather than SEQNO\_1, backstop, SEQNO\_N+1, etc...

At that point I made a test release and modified BDSPillAccessor to print the VldTimeStamp it was passed and to give a traceback. Access to BDSPillAccessor was coming from two places. The problem I've found is there is a fundamental clash between what AtNuOutput/NtpMaker::FillSpillInfo() is doing and what Filtration/DataQualityInterface::ProcessBeamStatus() is doing.

In DataQualityInterface it starts with a VldContext and uses SpillTimeFinder::GetNearestSpill() to construct a \_new\_ time with which it will query the BDSPillAccessor. But in NtpMaker, while it does get a SpillTimeND from SpillTimeFinder, it nevertheless uses the time in the VldContext when querying BDSPillAccessor.

Problems arise when SpillTimeFinder returns a new time that isn't in the same SEQNO of BEAMMONSPILL (presumably because the beam was actually off). Then the two start thrashing each other's DBI caches and each new record triggers a query for either (the real nearest spill) or (the BEAMMONSPILL that covers the record's VldContext).

I think the easiest solution, but Andy must say if this somehow invalidates some assumption, that AtNuOutput/NtpMaker.cxx lines 1310 to 1312 change from:

```
const SpillTimeND& ndspill = stf.GetNearestSpill(vldc);
```

```
const BeamMonSpill* bmspill = BDSpillAccessor::Get().LoadSpill(vldc.GetTimeStamp());
```

to:

```
const SpillTimeND& ndspill = stf.GetNearestSpill(vldc);
```

```
VldTimeStamp ts = vldc.GetTimeStamp();
```

```
if ( ndspill ) ts = ndspill.GetTimeStamp()
```

```
const BeamMonSpill* bmspill = BDSpillAccessor::Get().LoadSpill(ts);
```

---

## History

**05/25/2011 10:00 am - Art Kreymer**

- % Done changed from 30 to 60

Date: Wed, 25 May 2011 15:47:45 +0100 (BST)

From: Andy Blake <blake@hep.phy.cam.ac.uk>

To: Robert Hatcher <rhatcher@fnal.gov>

Cc: Minos Mailing List <MINOS\_SOFTWARE\_DISCUSSION@listserv.fnal.gov>

Subject: Re: MINOS: why "atmos" jobs beat up the database

Hi Robert,

Thanks very much indeed for finding the problem!

I've committed your suggested change to CVS. Running some test jobs, it looks to me like the repeated db queries go away.

Best Wishes,

Andy

**05/25/2011 10:00 am - Art Kreymer**

- % Done changed from 60 to 70

Date: Wed, 25 May 2011 09:52:34 -0500

From: Howard Rubin <rubin@iit.edu>

To: Andy Blake <blake@hep.phy.cam.ac.uk>

Cc: Robert Hatcher <rhatcher@fnal.gov>,

Minos Mailing List <MINOS\_SOFTWARE\_DISCUSSION@listserv.fnal.gov>

Subject: Re: MINOS: why "atmos" jobs beat up the database

Robert,

I've shut down the cron job and current processing should finish in somewhat under 2 hours.

Could you please then backport the changes to R2.0.5?

**05/25/2011 10:46 am - Art Kreymer**

Date: Wed, 25 May 2011 10:41:46 -0500

From: Robert Hatcher <rhatcher@fnal.gov>

Okay, backport onto R2.0 branch is done (and rebuilt); tag for R2.0.5 has been pushed; code has been checked out into R2.0.5 release. All I need is a go ahead and I can rebuild the R2.0.5 AtNuOutput package.

-robert

**05/25/2011 11:12 am - Art Kreymer**

Date: Wed, 25 May 2011 11:03:01 -0500

From: Rashid Mehdiyev <rmehdi@fnal.gov>

Hi Robert,

Please proceed with that rebuild.

Thanks,  
Rashid

**05/25/2011 02:11 pm - Art Kreymer**

Date: Wed, 25 May 2011 11:53:25 -0500

From: Robert Hatcher <rhatcher@fnal.gov>

Done. (Sorry about the delay... I was consulting with someone and not checking my mail).

**05/27/2011 07:22 am - Robert Hatcher**

- *Status changed from New to Resolved*

- *% Done changed from 70 to 100*

The fix has been implemented; backported to R2-0 branch; R2.0.5 tag for appropriate files has been pushed; the FNAL R2.0.5 (w/ and w/o -O) AtNuOutput has been re-built; running resumed. Jobs that were taking 50+ minutes now take less than half that.